AD-A100 654 GEORGIA INST OF TECH ATLANTA SCHOOL OF INFORMATION A--ETC F/G 9/2
RESEARCH PROGRAM IN FULLY DISTRIBUTED PROCESSING SYSTEMS.(U)
MAR 81 N00014-79-C-0873
UNCLASSIFIED NL

END
DATE FILMED
7-81
DTIC

LEVEL II

12

AD A100654

School of

Information and Computer Science

GEORGIA INSTITUTE OF TECHNOLOGY

THE GEORGIA INSTITUTE OF TECHNOLOGY

RESEARCH PROGRAM IN

FULLY DISTRIBUTED PROCESSING SYSTEMS.

Quarterly Progress Report Number 6
1 December 1980 - 28 February 1981.

March 1981

Supported by

Office of Naval Research (ONR)
Contract: N00014-79-C-0873
GIT Project: G36-643

U.S. Air Force Rome Air Development Center (RADC)
Contract: F30602-78-C-0129
GIT Project: G36-654

National Science Foundation (NSF)
Contract: MCS77-28305
Subcontract from Univ. of Wisc.: 144-L729
GIT Project: G36-630

U.S. Army Research Office (ARO)
Contract: DAAG29-79-C-0155
GIT Project: G36-638

U.S. Army Institute for Research in
Management Information and Computer Science (AIRMICS)
Contract: DAAK70-79-D-0087
GIT Project: G36-647

School of Information and Computer Science
Georgia Institute of Technology
Atlanta, Georgia 30332

1. INTRODUCTION

This is the Sixth Quarterly Progress Report prepared on the Georgia Tech Research Program in Fully Distributed Processing Systems (FDPS).

a. Program Description.

The Georgia Tech Research Program in Fully Distributed Processing Systems is a comprehensive investigation of data processing systems in which both the physical and logical components are extremely loosely coupled while operating with a high degree of control autonomy at the component level. The definition of the specific class of multiple computer systems being investigated, and the operational characteristics and features of those systems is motivated by the desire to advance the state-of-the-art for that class of systems that will deliver a high proportion of the benefits currently being claimed for distributed processing systems. The scope of individual topics being investigated under this program ranges from formal modeling and theoretical studies to empirical examinations of prototype systems and simulation models. Also included within the scope of the program are areas such as the utilization of FDPS's and their interaction with management operations and structure.

b. Program Support.

The principle support for the program is a Selected Research Opportunity contract from the Office of Naval Research; however, there are a number of other sources of funding which also support the program. A list of the currently active contracts and grants is given below.

Title: "Research on Fully Distributed Data Processing Systems"
Funding Agency: Office of Naval Research (ONR)
Contract Number: N00014-79-C-0873
GIT Project No.: G36-643
Principle Investigator: Philip H. Enslow, Jr.

Title: "Evaluation of Distributed Control Models"
Funding Agency: U.S. Air Force Rome Air Development Center (RADC)
Contract Number: F30602-78-C-0120
GIT Project No.: G36-654
Principle Investigator: Philip H. Enslow, Jr.

Title: "Foundations of Deterministic Scheduling of Processes for Parallel Execution"
Funding Agency: National Science Foundation (NSF)
Contract Number: MCS77-28305
(Univ. of Wisc. subcontract number: 144-L729)
GIT Project Number: G36-630
Principle Investigator: Richard A. DeMillo

Title: "Theory of Systems of Asynchronous Parallel Processors"
Funding Agency: U.S. Army Research Office (ARO)
Contract Number: DAAG29-79-C-0155
GIT Project Number: G36-638
Principle Investigator: Nancy Lynch

Title: "Support of MILPERCEN Data Storage Concept"
Funding Agency: U.S. Army Institute for Research in Management Information and Computer Science (AIRMICS)
Contract Number: DAAK70-79-D-0087
GIT Project Number: G36-647
Principle Investigator: A.P. Jensen

2. ORGANIZATION AND STAFFING

Faculty

Davida, George--Professor
DeMillo, Richard A.--Associate Professor
Enslow, Philip H. Jr.--Professor
Griffeth, Nancy--Assistant Professor
Jensen, Alton P.--Professor
LeBlanc, Richard--Assistant Professor
Livesey, Jon--Assistant Professor
Lynch, Nancy--Associate Professor
Miller, Raymond--Professor

Staff

McDonell, Sharon--Senior Secretary
Myers, Jeanette--Research Scientist
Pinion, Nancy--Part-time Secretary
Mongiovi, Roy--Research Technologist I

Students

There are approximately 30 students working on various projects in the FDPS Research Program. Of these, 12 are in the Ph.D. program and 5 are preparing their MS Thesis on topics in FDPS.

3. CURRENT RESEARCH PROJECTS

The specific research projects have been organized into the major areas identified in the basic program proposal.

A. Theoretical and Formal Studies

A.1 Models of Asynchronous Processors
A.2 Decomposition of Parallel Systems
A.3 Reliable Systems
A.4 Time Performance of Distributed Systems
A.5 Audit Algorithms
A.6 Ticket Systems

A.7 Synchronous Simulation
A.8 Distributed Resource Allocation
A.9 Theory of Distributed Databases
A.10 Arbiter Design
A.11 Shared Memory Bounds for Synchronization Problems
A.12 Mutual Exclusion
A.13 Adaptive Distributed Resource Allocation Algorithms
A.14 Using Complementary Distributed System Models
A.15 Probabilistic Algorithms in Distributed Systems
A.16 Stochastic Synchronization
A.17 Research Allocation in a Failure-Prone Environment

B. Physical Interconnection and Networking

B.1 Heterogeneous Networking
B.2 Local Networking in Fully Distributed Processing Systems

C. Distributed Operating Systems

C.1 Decentralized and Distributed Control
C.2 Resource Allocation and Work Distribution in an FDPS
C.3 Fully Distributed Operating System - Initial Considerations
C.4 Local Operating System
C.5 Communications Support for Distributed Systems
C.7 FDOS - Preliminary Implementation Studies

D. Distributed Data Bases

D.1 Concurrency Control in Distributed Database Systems
D.2 Support of MILPERCEN Data Storage Concept
D.3 Implementation of the Audit Algorithm

E. Fault-Tolerance

F. Special Hardware to Support FDPS

G. Application of Distributed Processing

H. System Design Methodologies

H.1 FDPS Requirements Engineering Techniques
H.2 Coordinating Large Programming Projects

I. System Utilization

I.1 A Language for Distributed Programming
I.2 System Implementation Language Development
I.3 Experiments with a Distributed Compiler

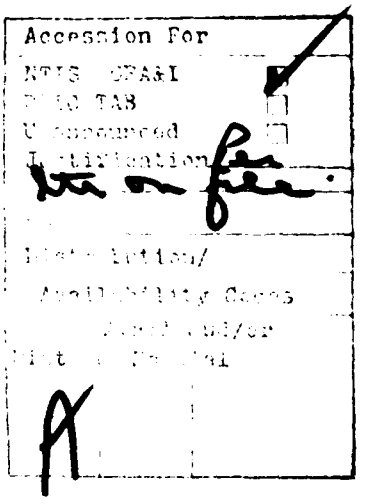

J. Security

J.1 Process Structures
J.2 System Security

K. System Management

L. Evaluation and Comparison

L.1 Simulation of Distributed Algorithms (Griffeth, Lynch)
L.2 Survivability

M. FDPS Testbed

M.1 Establishment of FDPS Testbed Facility
M.2 Remote Load Emulator
M.3 Fully Distributed Operating System Simulation Testbed

4. SUMMARY OF PROGRESS

A.1 Models of Asynchronous Processors (Lynch, Fischer)

No further progress anticipated. Project terminated.

A.2 Decomposition of Parallel Systems (Lynch, Fischer)

No significant progress to report.

A.3 Reliable Systems (Lynch, Fischer, Lamport)

No significant progress to report.

A.4 Time Performance of Distributed Systems (Lynch, Fischer, Lazowska, Schönhage)

No significant progress to report.

A.5 Audit Algorithms (Griffeth, Fischer, Lynch)

The earlier bank audit algorithm has been generalized considerably to an algorithm that returns a global state of a very general distributed system (e.g. a distributed data base), without halting concurrent operations in progress. The new general algorithm can be used for failure detection and recovery in distributed systems , and consistency checking in data base systems. It appears to be quite fast. The first draft of an invited paper, "Global States of a Distributed System", was prepared for presentation at the 1981 IEEE Conference on Distributed Software and Data Bases.

A.6 Ticket Systems (Fischer, Griffeth, Guibas, Lynch)

The paper, "Optimal Placement of Identical Resources in a Distributed Network", was rewritten for the Paris Conference on Distributed Systems.

A.7 Synchronous Simulation (Lynch, Fischer, Arjomandi)

A paper, "A Difference in Efficiency Between Synchronous and Asynchronous Systems", was written and submitted for presentation at the 1981 SIGACT Conference.

A.8 Distributed Resource Allocation (Lynch)

The paper, "Fast Allocation of Nearby Resources in a Distributed System", was revised for invited publication under the title, "Upper Bounds on Static Resource Allocation in a Distributed System, in a special issue of the Journal of Computer and System Sciences based on the 1980 SIGACT Conference.

A.9 Theory of Distributed Databases (Lynch, Griffeth)

The Audit Algorithm (Project A.5) has been generalized to apply to very general distributed data base systems. Discussions have been carried out attempting to generalize the usual notion of "serializability" used for correctness in data bases.

A.10 Arbiter Design (Lynch, Griffeth, Schönhage, Fischer)

No significant progress to report.

A.11 Shared Memory Bounds for Synchronization Problems (Burns, Lynch)

No further progress anticipated. Project terminated.

A.12 Mutual Exclusion (Burns, Lynch)

This Project was completed with the publication of a Ph.D. Thesis by J. Burns.

A.13 Adaptive Distributed Resource Allocation Algorithms (Ghoudjehbaklou, Lynch)

No further progress anticipated. Project terminated.

A.14 Using Complementary Distributed System Models (Lynch, Rounds, Miller)

No significant progress to report.

A.15 Probabilistic Algorithms in Distributed Systems (Lynch, Arjomandi, Fischer)

No significant progress to report.

A.16 Stochastic Synchronization (DeMillo, Miller, Lipton)

The work in Stochastic Synchronization has been brought to publication with an article, "Stochastic Synchronization". This paper reports results of simulations to support analytic results.

A.17 Resource Allocation in a Failure-Prone Environment (Fischer, Lynch, Burns, Borodin)

Revision of a paper, "*Resource Allocation with Immunity to Limited Process Failure*", is being carried out for journal publication.

B.1 Heterogeneous Networking

Initial task completed. No further work in this area during this quarter. Project terminated.

B.2 Local Networking in FDPSs (Enslow)

Work during this period has been focused on the control and software problems of the local network that has been installed. Emphasis is being placed on software and system reliability as well as the incorporation of important features such as a network name server.

C.1 Decentralized and Distributed Control (Enslow, LeBlanc, Saponas)

The first task under this project, identification and definition of models of distributed control, was completed with the publication of the final report. Work on the second task, evaluation of these models, is proceeding.

C.2 Resource Allocation and Work Distribution in an FDPS (Enslow, Sharp)

A draft of a paper entitled, "An Initial Examination of Resource Management and Work Distribution in a Fully Distributed Processing System", was completed and is expected to be published as a technical report in the near future. A draft version of a proposal for a simulation experiment to measure the performance of work distribution algorithms in fully distributed systems was completed. The proposal will also be published as a technical report.

C.3 FDOS - Initial Considerations (Enslow, Livesey, LeBlanc, Akin, Flinn, Forsyth, Fukuoka, Maccabe, Myers, Pitts, Saponas, Skowbo, Spafford)

Many of the initial considerations for the design and implementation of a Fully Distributed Operating System have appeared under other projects, primarily C.1, C.2, and C.4. The local operating system, Project C.4, has been designed to support alternative implementations of distributed operating system concepts. Project C.3, as well as Project C.7, will be terminated this quarter with the conclusion that our efforts are best directed towards the implementation of a local operating system that will provide a real testing environment for distributed operating system concepts.

C.4 Local Operating System (Livesey, LeBlanc, Spafford, Myers, Flinn, Forsyth, Fox, Fukuoka, Greene, Hopkins, Mongiovi, Pitts)

A preliminary design is complete and we are in the process of expanding this design, top-down, to the coding level. In anticipation of implementing an LOS on the PR1MEs, a special class, ICS 8113-F, "The Organization, Architecture, and Programming of PR1ME Computers", was offered this quarter by a member of the Local Operating System team. Also, in conjunction with this project, ICS 8113-L, a seminar on distributed operating system concepts, was offered.

C.5 Communications Support for Distributed Systems (Enslow, Skowbo, Wice)

At the commencement of the Sixth Quarter, three subtasks for this project had been identified:

1. Development and evaluation of a highly distributed routing algorithm for a message-switching network.
2. Completion of an enhanced Interprocess Communication Simulator to support communication research, particularly the evaluation phase of the third subtask.
3. Development and evaluation of a highly selective acknowledgement protocol for error control in fully distributed systems and other computer networks requiring improved throughput on high-speed, long-delay satellite links.

The first subtask has been abandoned for several reasons:

1. The physically tight coupling of the current hardware configuration in the Computer Lab would not permit a direct evaluation of routing methods, since communication between processors is constrained to follow a single fixed path, and there are no plans to change this configuration in the near future.
2. The development of software for a completely operational communication subsystem or its simulation is well beyond the limitations of this project.
3. Details of the proposed implementation reveal unforeseen operational difficulties which would largely negate the expected benefits of this algorithm, particularly in comparison with currently available alternatives which are more centralized and tightly-coupled, but highly adaptive and refined by years of operational experience.

The second subtask is very near completion, however, its exact status cannot be determined, pending a demonstration, evaluation, and final report in the form of a Masters Thesis by Mr. Wice. (Mr. Wice has left Georgia Tech for full-time employment. He anticipates completion of his thesis during the next quarter.)

The highly selective acknowledgement protocol has been described in some detail with a plan for undertaking its evaluation; however, further progress on this third subtask has been deferred, pending a more thorough study of the communications support required by fully distributed processing systems.

C.7 FDOS - Preliminary Implementation Studies (Myers, Enslow)

This project has been terminated as described above in the progress report on Project C.3.

D.1 Concurrency Control in Distributed Database Systems (Griffeth, Livesey, Lynch)

A model of a distributed database system has been developed for use in simulating distributed concurrency control algorithms. Design of the simulation is in progress.

D.2 Support of MILPERCEN Data Storage Concept (Jensen, Doyle, Gehl, Bingham)

This quarter's work included typing of the transcript of the workshop, "Implications of Data Base Technology for Human Resource Information Management", held in October, 1980. The transcript was sent to the workshop attendees for review and for approval concerning the basic ideas put forth at the workshop. Secondly, initial work was begun in reviewing the literature of data base technology from the perspectives of design, performance, data communications, and evaluation of data bases. Third, the project staff was invited on January 16, 1981 by AIRMICS to an IBM General Systems presentation on data base standards, IMS products and systems, and the data base environment. Fourth, the writing of the final project report was begun and is continuing.

D.3 Implementation of the Audit Algorithm (Griffeth, Livesey, Lynch)

An investigation of block allocation by the Series/1 file management system is in progress, for a test implementation of the audit algorithm.

H.1 FDPS Requirements Engineering Techniques (Underwood, Corley)

No further progress anticipated. Project terminated.

H.2 Coordinating Large Programming Projects (Enslow, Smith)

The initial formulation of the model of the communcation process in the development of large software systems has been completed. The model is now undergoing refinement and amplification.

I.1 A Language for Distributed Programming (LeBlanc, Maccabe, Hardin)

Design work has continued during this quarter and some implementation work has started. A paper was presented at the Louisiana Computer Exposition, providing some useful interaction with other researchers in this area.

I.2 System Implementation Language Development (LeBlanc, Akin)

The code generator has been completely designed. Implementation will start in the near future.

I.3 Experiments with a Distributed Compiler (LeBlanc, Moore)

Moore is currently writing his M.S. thesis describing this work.

J.1 Process Structures (DeMillo, Lipton, Miller)

We have been preparing a book-length treatment of cryptographic protocols for publication in late 1981. As part of this project a large number of protocols suited to distributed systems have been identified and potential lines of compromise explored. We have also identified several new algorithms for examination. Theoretical research has centered on quantitative measures of system security .

J.2 System Security (Livesey, Davida, DeMillo)

Operating systems security is a relatively new area. Harrison, Ruzzo, and Ullman have shown that the safety question for operating systems is undecidable. However, practical design issues require that new approaches to secure operating systems be developed. Davida, DeMillo, and Lipton have introduced a new architecture that implements the "star" property for multilevel secure operating systems. The approach differs from that of other designs which rely on verification techniques to implement a secure kernel.

A paper is being written in which we present a new architecture that achieves security in a timeshared operating system. This is a very important class of operating systems since they are widely used.

L.1 Simulation of Distributed Algorithms (Griffeth, Lynch)

The ticket simulation program has been transferred to the PRIME and is being tested there. A graphical display of the simulated allocation of tickets has been developed.

L.2 Survivability (DeMillo, Martin)

This Project was completed with the publication of a Ph.D. Thesis by E. Martin on experimental aspects of survivability in distributed systems. A central result of this thesis was a factor analysis of approximately 300,000 data points to identify key parameters which influence system survivability.

M.1 Establishment of FDPS Testbed Facility (Myers, Fox)

A high-level language interface to PR1MENET's X.25 subroutines has been implemented. This interface allows asynchronously running programs to communicate using send and receive primitives.

M.2 Remote Load Emulator (Myers, Enslow, Forsyth)

No significant progress to report.

M.3 FDOS Simulation Testbed (LeBlanc, Hopkins, Myers)

The simulator is near completion. Some tailoring is being done to accommodate Project C.2, which will soon require this simulator.

5. TRAVEL RELATED TO THE FDPS PROGRAM

Date of Trip: 15-17 December, 1980
Individual(s) Traveling: Richard LeBlanc & Nancy Lynch
Itinerary: Fallbrook, California
Purpose: Attend a workshop on fundamental issues in distributed computing

Date of Trip: 3-5 January, 1981
Individual(s) Traveling: Richard DeMillo
Itinerary: San Francisco, California
Purpose: Present invited talk at annual meeting of American Mathematical Society

Date of Trip: 9-10 February, 1981
Individual(s) Traveling: Philip H. Enslow, Jr.
Itinerary: Virginia Polytechnic Institute and State University, Blacksburg, Virginia
Contact: Roger Ehrich
Purpose: Present talk on FDPS Program to faculty and students

Date of Trip: 26-27 February, 1981
Individual(s) Traveling: Richard LeBlanc
Itinerary: Lafayette, Louisiana
Purpose: Present a paper at the Louisiana Computer Exposition

6. VISITORS

Dates of Visit: 10-13 January, 1981
Visitor: Mike Fischer
Contact: Nancy Lynch & Nancy Griffeth
Purpose: Work on Project A.5 and participate in the Ph.D. Dissertation Defense of J. Burns.

7. PUBLICATIONS

Author(s): N. Lynch
Title: Fast Allocation of Nearby Resources in a Distributed System
Type: Conference paper
Status: Published
Publ. Date: May, 1980

Author(s): P.H. Enslow, Jr.
Title: Quarterly Progress Report - Number 5
Type: Quarterly Progress Report
Status: Published
Publ. Date: December, 1980

Author(s): E.W. Martin
Title: Survivability in Gracefully Degrading Computer Systems.
Type: Ph.D. Thesis
Status: Published
Publ. Date: January, 1981

Author(s): R.A. DeMillo
Title: Cryptographic Protocols
Type: Conference paper
Status: Published
Publ. Date: January, 1981

Author(s): J. Burns
Title: Complexity of Communication among Asynchronous Parallel Processes
Type: Ph.D. Thesis
Status: Published
Publ. Date: 12 January, 1981

Author(s): P.H. Enslow, Jr. & T.G. Saponas
Title: Distributed and Decentralized Control in Fully Distributed Processing Systems - A Survey of Applicable Models
Type: Final Technical Report
GIT Number: GIT-ICS-81/02
Status: Published
Publ. Date: February, 1981

Author(s): R.J. LeBlanc
Title: Communication and Control Abstractions in a Programming Language for Fully Distributed Systems
Type: Conference paper
Status: Published in Proceedings of the Third Annual Louisiana Computer Exposition
Publ. Date: February, 1981

Author(s): R.J. LeBlanc & A.B. Maccabe
Title: P+D: Language Features for Distributed Programming
Type: Technical Report
Status: In preparation
Publ. Date: March, 1981

Author(s): R.A. DeMillo, R. Lipton, & R. Miller
Title: Stochastic Synchronization
Type: Conference paper
Status: Submitted for publication
Publ. Date: March, 1981

Author(s): G. Davida, R. DeMillo, & R. Lipton
Title: Achieving Multilevel Security Through Distributed Systems
Type: Conference paper
Status: Submitted for publication
Publ. Date: April, 1981

Author(s): M. Fischer, N. Griffeth, L. Guibas, & N. Lynch
Title: Optimal Placement of Identical Resources in a Distributed Network
Type: Conference paper
Status: Accepted by Paris Conference on Distributed Systems
Publ. Date: April, 1981

Author(s): E. Arjomandi, M. Fischer, & N. Lynch
Title: A Difference in Efficiency Between Synchronous and Asynchronous Systems
Type: Conference paper
Status: Accepted by 1981 SIGACT Conference
Publ. Date: May, 1981

Author(s): M. Fischer, N. Griffeth, & N. Lynch
Title: Global States of a Distributed System
Type: Conference paper
Status: Invited by 1981 IEEE Conference on Distributed Software and Data Bases
Publ. Date: July, 1981

Author(s): N. Lynch
Title: Upper Bounds on Static Resource Allocation in a Distributed System
Type: Journal paper
Status: Invited, revised, and submitted to Journal of Computer and System Sciences

Author(s): M. Fischer, N. Lynch, J. Burns, & A. Borodin
Title: Resource Allocation with Immunity to Limited Process Failure
Type: Conference paper
Status: Revision in preparation